An Ideal Navigator Model of Human Wayfinding: Learning One's Way Around a New Town

Jeremy Manning[†‡], Michael J. Kahana[§] & Robert Sekuler[†]
[†]Volen Center for Complex Systems, Brandeis University
[‡]Neuroscience Graduate Group, University of Pennsylvania
[§]Department of Psychology, University of Pennsylvania

Draft from October 5, 2006

Searching for a target in an unfamiliar environment requires acquiring, storing, processing, and recalling spatial information. We use datasets from two virtual reality navigation experiments to design and validate a simple biologically-inspired computational model of these processes. Subjects in both datasets play taxicab drivers, picking up and delivering passengers in a series of small virtual towns. After just a few deliveries, subjects learn to generate minimum-distance, novel paths between passengers and their destinations. Subjects' rapid learning and near-ceiling performance inspires the creation of an ideal navigator model, which makes optimal use of sensory information to learn navigationally relevant spatial information about the environment as quickly as possible. We use MAGELLAN, a simple, two-parameter, ideal navigator model, as a benchmark against which to assess possible sources of subject error. By systematically degrading the ideal navigator's vision and memory, MAGELLAN accounts for human subjects' mean performance in both datasets, and correctly predicts the difficulty that subjects encounter in navigating different environments.

Goal-directed navigation is among the most important behaviors in the repertoire of any non-sessile animal, humans included (Montello, 2005). Human goal-oriented navigation has long been the subject of empirical and theoretical study, but acquired new impetus from two recent developments. The first is the discovery in the mammalian brain of neurons whose characteristics seem well suited to encoding information that is relevant to navigation; the second is the availability of virtual reality technology capable of producing appropriate testbeds for studying navigation. We first briefly review both of these developments. We then introduce an ideal navigator model and use that model to help explain data from two experiments in which subjects learned to navigate computer-generated towns. Finally, we generalize our model by degrading its vision and memory. This degraded ideal navigator is able to account for variation in the difficulty of learning different environments, as well as experience-dependent improvements in navigation performance.

Navigation-relevant neurons

In recent years, researchers have identified several classes of neurons that participate in navigation (Muller, Stead, & Pach, 1996). Such neurons represent sensory and mnemonic information that could be exploited in the generation of a cognitive map of the environment (Tolman, 1948; Sholl, 1987; McNaughton, Battaglia, Jensen, Moser, & Moser, 2006). The most well-known of these navigation-relevant neurons are hippocampal *place cells*, neurons that are responsive to specific locations within an environment irrespective of the animal's viewpoint (Muller et al., 1996; O'Keefe, 1979; Ekstrom et al., 2003), and postsubicular *head-direction cells*, neurons that are responsive an animal's headings within the environment (Robertson, Rolls, Georges-Francois, & Panzeri, 1999; Taube, Muller, & Ranck, 1990). Equally complex, intriguing properties define parahippocampal *view cells*, which respond when specific landmarks are seen (Rolls & O'Mara, 1995; Ekstrom et al., 2003), and *goal-responsive cells*, which are activated when a behaviorally-relevant goal becomes visible (Ekstrom et al., 2003; Hok, Save, Lenck-Santini, & Poucet, 2005; Hollup, Molden, Donnett, Moser, & Moser, 2001).

Significant progress is being made in understanding the ways in which environmental features and feature locations are represented in the brains of navigating animals. However, we are far from having a complete account of human navigation. We do not know how the representations of individual locations and goals are translated into action, particularly into action that takes an animal along a novel, but efficient path between locations; we do not know how the representations might be organized into a hierarchy such as might be needed for the "center and spoke" arrangement that characterizes many environments (Voicu, 2003; Graham, Joshi, & Pizlo, 2000); and we do not know how time constraints or other subsidiary goals, such as path minimization (Vieira,

Supported by NIH grants MH61975, MH55687, and MH068404, and by NSF grants SBE-0354378 and BCS-0527689. We are grateful for the assistance of Josh Jacobs, Igor Korolev, Matt Mollison, and Ben Okaty.

Doria Neto, & Costa, 2003), might be enacted neurally. Even less well understood is the ability to predict cooperative behaviors among individual navigators (Qi & Sun, 2005).

Virtual reality as an experimental paradigm

Over the last decade, virtual reality (VR) technology has made possible a range of applications, including walkthrough-previews for prospective buyers of real estate properties, training for the military and first responders, desensitization therapies for phobias, regimens for rehabilitation after brain injury, and video games on which a multi-billion dollar industry rests. Technological advances make it possible to generate customized, interactive VR environments. The persuasive realism of interactive environments (Sanchez-Vives & Slater, 2005), and the degree to which interactions can be manipulated, measured, and analyzed, makes them well-suited for empirical research.

For empirical studies of navigation, VR is particularly valuable when it allows researchers to build multiple environments from a common set of generative rules, and to compare performance across those environments, as we have done in our research. For example, by using explicit rules to govern the layout of multiple environments and/or the placement of structures within those environments, researchers can identify the principles that govern navigation and learning, independent of the influence of any single environment's singularities (Newman et al., 2006). In this way, results from studies of wayfinding in properly designed VR environments are arguably more generalizable than results from any single real world environment.

Of course, even the most compelling VR is not likely to be confused with the real world. However, so long as a VR environment has sufficient realism, it can be a useful testbed within which to study wayfinding. With sufficient realism, environmental knowledge gained in a VR setting can transfer to the real world (Witmer, Baily, Knerr, & Parsons, 1996). The very same VR towns that were used to generate the datasets analyzed in this paper were sufficiently realistic to activate navigationally-useful place, view, and goal cells in the brains of humans who navigated through those towns (Ekstrom et al., 2003). Of course it is difficult to know precisely *a priori* how much realism is sufficient, or exactly what attributes *must* be represented in the environment. However, experimental tests of hypotheses can provide useful clues. For instance, we know that navigation performance degrades when the illusion of natural movement is removed, for example by reducing the rate of optic flow a subject experiences in a VR environment (Kirschen, Kahana, Sekuler, & Burack, 2000). We also know that distinctive building textures and sizes are important features of a VR environment. In fact, when an environment's texture is overly homogeneous, subjects frequently become lost, and tend to fall back upon simple list learning strategies, such as "turn left, then right turn, etc.," rather than learning the environment's spatial layout (Kirschen et al., 2000).

Preview of this report

Our principal goal is to develop a model for navigation learning, using virtual environments as testbeds. We present entirely new analyses of two previously reported studies (Newman et al., 2006; Korolev, Jacobs, Mollison, & Kahana, 2005) in which subjects functioned as taxicab drivers in virtual towns, picking up passengers and delivering them to target locations. In our analysis of these experiments, we compared human subjects' performance to that of MAGELLAN, a theoretical ideal navigator. As expected, MAGELLAN outperformed both experiments' human navigators. Then, by systematically degrading each of MAGELLAN's parameters, we attempted to identify the variables that kept human navigators from matching the performance of the ideal. When MAGELLAN's vision and memory were degraded sufficiently, the model's performance matched that of human subjects. In addition, the degraded ideal navigator successfully predicted the difficulty posed by various randomly-generated environments.

Empirical studies of navigation in VR environments

In both VR experiments described briefly below, subjects operated as taxicab drivers in small virtual towns, searching for randomly placed passengers and then taking each passenger to a requested destination, that is, a store. Over successive passenger searches, pickups, and deliveries, subjects learned to take passengers to their destinations along the shortest possible paths. Previously, the same paradigm has also been used in a series of electrophysiological investigations of human navigation (e.g., Caplan et al., 2003; Ekstrom et al., 2003, 2005; Kahana, 2006).

Table 1 summarizes the key differences between the two experiments. Although the general setup remained constant across both experiments, our model did seem to overcome small task-related differences between experiments, that is, changes in controller type, rate of movement, and size of the virtual environment.

Dataset One (Newman et al., 2006)

The first set of empirical results we analyzed were reported as the second experiment in Newman et al. (2006), where full details can be found. Whereas Newman et al. (2006)'s focus was on transfer of spatial knowledge across different environments, our own analysis of their data focused on subjects' ability to learn efficient paths through a given novel environment.

Methods. One hundred twenty-three Brandeis University undergraduates, with approximately equal numbers of males and females, participated for monetary compensation, which was supplemented with a performance-based bonus.

Each subject learned to navigate virtual towns that were laid out on a 5×5 orthogonal grid of streets, with one uniquely-textured *structure* centered on each block. The structures comprised two categories: twenty-one multistory

Table 1
Key Conditions Used in Generating Datasets One and Two

Condition	**Dataset One**	**Dataset Two**
Number of subjects tested	123	21
Environment size	5×5 blocks	6×6 blocks
Number of environments per subject	1	6
Stores in each environment	4	5
Deliveries per environment	20	15
Spatial arrangement of stores	Quasi-random	Constrained
Deliveries to each store	5	3
First delivery starts with passenger in taxi?	No	Yes
Controller	Arrow keys ($\uparrow, \downarrow, \leftarrow, \rightarrow$)	Joystick
Forward speed	1.17 $units/sec$	0-1 $unit/sec$
Turning speed	$20°/sec$	$50°/sec$
Field of View	56°	60°
Screen resolution	640×480	800 × 600
Screen refresh interval	30 *msec*	20 *msec*

office *buildings* and four one-story *stores*. Passengers were delivered only to stores. The office buildings varied slightly in width and in height, and were surrounded by lawn. The stores, which were all of the same size, had smaller footprints than the office buildings, and were surrounded by pavement. The outer boundary of the town was marked by a texture-mapped stone wall; no other visual information was available beyond the boundary.

After a familiarization session, each subject picked up and delivered twenty passengers, five to each of four stores. The deliveries were constrained such that passengers had to be delivered to each of the four stores before any store was delivered to a second time. In each environment, stores were distributed pseudo-randomly, subject to some weak constraints: all four stores could not simultaneously occupy the four corners of the town, and no two stores could occupy immediately adjacent blocks. Potential passengers were placed pseudo-randomly in the town, subject to the constraint that no passenger could be in the direct line of sight from the preceding target store. The randomization of passenger pickup locations served to minimize the degree to which memorization of individual routes, that is, memorizing one path for each starting and ending location, would be effective. Instead, subjects were forced to generate novel paths from unique passenger pickup locations to fixed targets (stores).

Each delivery began with a *passenger search*, during which the subject sought the next passenger. When the passenger was picked up, a text screen notified the subject of the next target store. The *delivery path* began with picking up the passenger and ended at the *goal* (the passenger's destination store).

The subject's virtual earnings as a taxi driver were continuously displayed in the upper right corner of the screen. The upper left corner of the screen contained a short description of the current instructions (e.g., "Find a passenger" or "Find the Java Zone"). Figure 1(a) is a screen capture taken during a passenger search.

Dataset Two (Korolev et al., 2005)

The second set of empirical results we analyzed were reported in Korolev et al. (2005), where full details can be found. Whereas Korolev et al. (2005)'s focus was on the electroencephalographic activity that accompanied navigation through various virtual environments, our analysis of those data focused on subjects' ability to learn efficient paths through a given novel environment. In addition, we supplemented the original dataset of fourteen subjects by testing seven additional subjects under identical conditions (Manning, 2006).

Methods. Twenty-one subjects were tested; fourteen at the University of Pennsylvania and seven at Brandeis University. All subjects (twelve male, nine female) participated for monetary compensation, with a performance-based bonus.

To make the navigation task more challenging, the virtual towns were laid out on a 6×6 block orthogonal grid of streets, rather than the 5×5 block grid used for Dataset One. In addition, the paradigm underwent some cosmetic changes (see Figure 1(b)) designed to enhance the realism of the virtual towns. These changes included an increase in the resolution with which images were displayed on the screen. Over the course of three sessions, subjects delivered passengers to five target stores in each of six unique VR towns.

Of the thirty-six blocks in each town, thirty-one were occupied by multi-story office buildings, and five held one-story stores. Passengers were delivered only to stores. The outer boundary of the town was marked by a texture-mapped brick wall; no other visual information was available beyond the boundary.

As in Dataset One's experiment, the deliveries were block-randomized. Passengers were scattered pseudo-

(a) Screen capture for Dataset One

(b) Screen capture for Dataset Two

Figure 1. Screen captures taken during a typical passenger search in experiments for Dataset One (upper panel) and Dataset Two (lower panel). In the upper panel, the Autopac store and several buildings are visible. In lower panel, the Butcher Shop store, a passenger, and several buildings are visible.

randomly throughout the environment, subject to the constraint that no passenger could be located in the direct line of sight from the preceding store.

After a familiarization session, subjects made deliveries in two unique environments per session, for a total of six unique VR towns by the end of their third session. In addition to maintaining store placement constraints from Dataset One's experiment, stores in Dataset Two's experiment were subjected to two additional constraints. The first additional constraint prevented pairs of stores from occupying blocks that were a "knight's move" away from one another; the second added constraint applied when stores shared either the same north-south or east-west alignment. Under that condition, the stores' locations had to be separated by more than two blocks along the orthogonal axis. These added constraints spread out the locations of stores, and prevented stores from lying within viewing distance of one another. The overall intent of the added constraints was to render various environments more uniform in their difficulty to be learned.

Also, unlike the experiment that generated Dataset One, here, for the first delivery in each new environment, a passenger was placed a short distance ahead of the subject, in plain view. This eliminated the first passenger search, during which subjects in Dataset One demonstrated rapid learning (Newman et al., 2006). As in Dataset One, though, successive deliveries began with a *passenger search*, during which the subject searched for the next passenger. When a passenger was picked up, a text screen notified the subject of the next target store. The *delivery path* began with picking up the passenger and ended at the *goal* (the passenger's destination store).

The upper left corner of the screen contained a short description of the current instructions (e.g., "Find a passenger" or "Find the Coffee Store"). Figure 1(b) is a screen capture from a typical passenger search.

Spatial learning

Across multiple deliveries, subjects learned to find nearly optimal paths from randomly placed passengers to fixed location stores. To measure the improvement in navigation over successive deliveries (i.e., the *spatial learning curve*) we subtracted the *ideal path distance* from the actual distance traveled between the passenger and the goal location. When a direct path is possible, the ideal path distance is equal to the *city block distance* ($\Delta x + \Delta y$) from the passenger to the goal store. When a physical obstacle or barrier would have blocked a direct path to the goal store (see Figure 2), we adjusted the ideal path distance to include the smallest detour that would avoid the obstacle. We refer to the difference between adjusted ideal path length and actual path length as *excess path length*. Note that by incorporating this adjustment for obstacles, the measure of excess path length is a realistic measure of the minimum delivery path distance that subjects could actually achieve.

As Figure 3 shows, in both datasets excess path length decreased rapidly over successive deliveries. Although the curves are quite similar in form, excess path length for the first delivery is larger in Dataset Two than in Dataset One. This is likely due to the methodological differences between the two studies, several of which were intended to increase the difficulty of learning that would be reflected in Dataset Two (see Table 1). Among these differences, the one that is likely to have accounted for the largest increase in first trial excess path length is Dataset Two's lack of a passenger search phase. Analyses of Dataset One revealed that subjects spent almost twice as long searching for the first passenger as they did delivering that passenger to his or her destination. This appreciable amount of exploration enabled subjects to learn a great deal about the environment and store locations prior to the very first delivery (Newman et al., 2006).

Dataset Two's subjects made deliveries in six unique en-

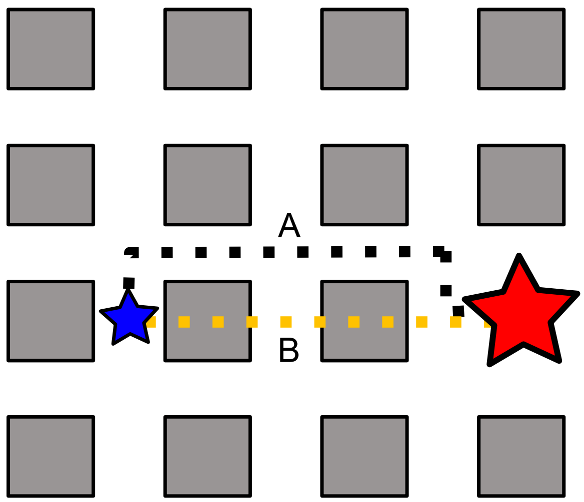

Figure 2. Obstacle avoidance. When the start position of the passenger (⋆) is directly aligned with the goal (⋆) in either the horizontal or vertical direction (but is not directly adjacent to the goal) the driver must navigate around one or more landmarks in order to make the delivery (curve A, shown in black). In such cases, the ideal path will have an excess path length equal to twice the city block distance from the passenger to the nearest intersection (the ideal path distance must account for a swerve in the beginning of the path and also at the end). Curve B, shown in orange, depicts the minimum distance path between the passenger and goal without taking environmental obstacles into account.

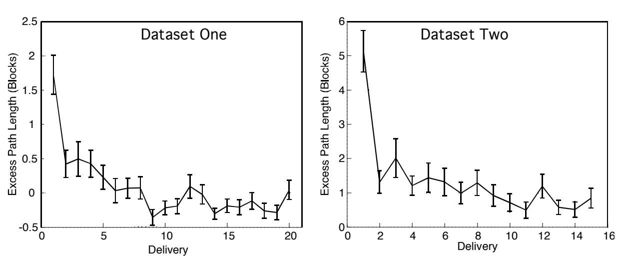

Figure 3. Mean excess path length in Dataset One (left panel) and Dataset Two (right panel). Subjects improve with the number of deliveries made. Note that a negative excess path length indicates that the ideal path distance (equal to the city block distance between the passenger and goal) is greater than the delivery path distance (the summed Euclidean distance between adjacent points along the subject's delivery path). Error bars represent ±1 SEM

vironments (two environments were learned successively on each of three days). Because subjects were exposed to multiple environments, we worried that the results might be influenced by warm-up or learning-to-learn effects (e.g., Keppel, Postman, & Zavortnik, 1968). To test this, we compared performance on the six successive environments that any subject encountered. A one-way ANOVA (analysis of variance) indicated that the population means of excess path length in each environment were unchanged over successive environments ($F(5,20) = 0.23, p > 0.95$). The clear stability in performance over successive, novel environments suggests that subjects were learning environment-specific features, such as the locations of particular stores, rather than some general navigation strategy or the rules that governed placement of target stores. The consistency in performance across six unique environments suggests also that subjects' exploration strategies remained relatively constant, even as subjects gained more experience with the task. Accordingly, we felt justified treating each unique combination of an environment with a subject as a distinct entity for the purposes of modeling and analysis.

MAGELLAN: a model of human spatial navigation

Both of the datasets we analyzed demonstrated rapid learning and near-perfect performance after just a few deliveries. We propose that the acquisition of ability to generate novel, efficient paths between arbitrary locations in a new environment implicates a spatial memory or cognitive map (Tolman, 1948) that is built up as the subject explores, and a route-generation mechanism that produces efficient paths between locations stored in that spatial memory. Neural systems that could navigate using a spatial map have been proposed by several researchers (for example, Gerstner & Abbott, 1997; Byrne, Becker, & Burgess, in press; Samsonovich & McNaughton, 1997).

Our proposed model of human spatial navigation is called MAGELLAN, whose eponym is the Portuguese explorer Ferdinand Magellan, commander of the first ship to circumnavigate the globe. MAGELLAN consists of the three modules shown in Figure 4. *Vision* provides a means for acquiring new information about environments. A *cognitive map* stores acquired spatial information. A *route generation* mechanism computes an efficient route between locations stored in the spatial memory. The interactions among modules provide a means for learning the spatial layout of a new environment and searching for unknown targets in the environment.

First, we define an ideal navigator version of MAGELLAN. In this special case of the general MAGELLAN model, each of the three modules functions optimally, that is, without noise. In more general variants of MAGELLAN any of the modules could be systematically degraded in an effort to approximate human-level performance. To anticipate our results, with vision and spatial memory modules degraded to appropriate degrees, MAGELLAN's performance matches that of human subjects in Datasets One and Two. We chose not to degrade MAGELLAN's third module, the route generation module. In the relatively small environments used for both datasets, the assumption of an optimal route generation mechanism seems to be appropriate.

The ideal navigator as benchmark

> *"If we could first know where we are, and whither we are tending, we could then better*

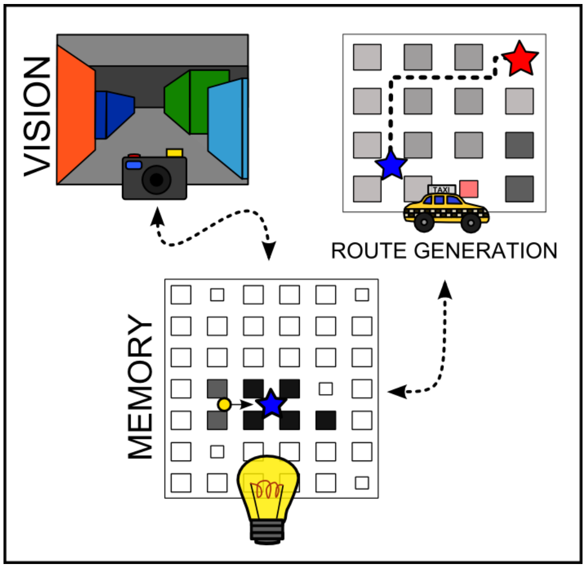

Figure 4. Modules and interactions in MAGELLAN. The *Vision* module provides a means for acquiring new information about the photic environment. The *Memory* module stores acquired spatial information. In the diagram, stores are represented in red, and office buildings are represented in black. Darker colored squares corresponds to a "stronger memory" of a particular structure, while white squares represent locations whose structures are not encoded in the spatial memory. A *Route Generation* module computes an efficient route between locations stored in the spatial memory. When the passenger's target destination is not stored in the spatial memory, the navigator's current location (shown by the yellow disc) is compared to "blank" blocks on the cognitive map (shown in white). The navigator uses the route generation module to generate an efficient path to the nearest unknown block. As the path is navigated, the vision module feeds new data to the spatial memory. In the degraded ideal navigator, the *Vision* and *Memory* modules are degraded, while the *Route Generation* module is left intact.

judge what to do, and how to do it."

– Abraham Lincoln, 1858 (Springfield, IL)

As Abraham Lincoln recognized in his "House Divided" address at the 1858 Republican State Convention, efficient navigation to a goal requires knowing one's current state ("where we are") and the desired state ("whither we are tending"). By definition, an ideal navigator armed with knowledge of its current position in an environment and the position of an intended destination (target), will generate an *ideal* path from its current position to the target. An ideal path, or an *optimal* path, is a delivery path whose length is equal to the minimal achievable path distance, taking into account the need to detour around impenetrable obstacles such as buildings. In the context of our task, a navigator is ideal when it (a) sees all structures (office buildings and stores) displayed on the computer screen, (b) encodes their spatial location without noise or forgetting, and (c) uses this information, in conjunction with knowledge of its own current location and heading, to generate ideal delivery paths to previously encoded goal stores. Note that an ideal navigator is not prescient: it cannot generate ideal paths to targets it has not yet seen. For goals that have not yet been seen, the ideal navigator (d) employs an efficient goal-search algorithm in order to quickly search the nearest, unknown sections of the town. Once the goal is seen or located on the subject's spatial map, the navigator's route-generation module produces the minimal achievable path from the navigator's current location to the goal.

The ideal navigator algorithm. From the pickup of a passenger to the delivery of that passenger to their desired target store, the route generation algorithm operates via a pair of instructions:[1]

1. **Find a goal.** If the location of the passenger's intended destination (the store) is already in spatial memory, go there. Otherwise, the goal is the closest block not yet associated with an office building or a store.[2] If there are multiple, equally-attractive possible goals, one is selected randomly.
2. **Take a step toward the goal.** If there are multiple, equally-efficient directions in which to step, one is selected at random.

We define one *step* by MAGELLAN as a move from one intersection in the town's roadways to an adjacent intersection. Consequently, the ideal navigator's heading will always be a multiple of 90°. This simplification allows us to focus on path distance—rather than the exact *shape* of a particular delivery path—as an index of spatial knowledge of the environment. Constraining the model's movement in this way ensures that unlike human subjects the model can never outperform the ideal path Although the ideal path length is computed using a city block distance adjusted for impenetrable obstacles, the human subject's delivery path could cut corners at headings other than 90°. Thus, the true ideal subject path is actually somewhere between the city block distance and the Euclidean distance from passenger to goal. In order to make a fair performance comparison between subjects and the model, we simplified the subjects' paths to approximate the shape of path that would be taken by the model. Appendix I describes the method used to simplify subject paths for comparison with MAGELLAN's paths.

When MAGELLAN enters a novel environment, its spatial memory is a *tabula rasa*. Then, as it moves through the town, the model identifies structures that would have been visible from its position in the virtual town. Each visible office building and store is added at the appropriate location to the ideal navigator's spatial memory of the town. The spatial memory includes the spatial location and a unique identifier for each structure that had been visible, and fills up as more office buildings and stores are seen.

[1] Appendix II provides pseudocode for our full model.

[2] Recall that passengers are delivered only to stores, never to office buildings.

When searching for a goal store that is not represented in memory, the ideal navigator minimizes effort per delivery (Robles-De-La-Torre & Sekuler, 2004), exploring first the most proximate location(s) at which office buildings or stores have not yet been seen. An indirect consequence of this goal-seeking strategy is that previously traversed areas of the town will tend to be avoided while searching for the goal. This strategy is analogous to *inhibition of return*, a phenomenon that has been observed in visual search and other tasks (Klein, 2000; Bennett & Pratt, 2001). Inhibition of return manifests itself by increased tendency to orient to or attend to novel locations, which could promote efficiency in exploring unfamiliar environments or when foraging for food.

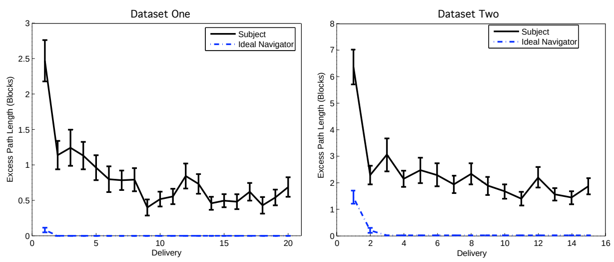

Figure 5. Ideal Navigator Predictions for Dataset One (left panel) and Dataset Two (right panel). Error bars represent ± 1 SEM. Recall that subjects in the experiment for Dataset One are given time to explore while looking for their first passenger, but subjects in the experiment for Dataset Two already have a passenger at the start of their session.

Ideal navigator vs. experimental data. We compared our ideal navigator's performance to subjects in each experiment, as shown in Figure 5. In the simulations, our goal was to determine whether the ideal navigator could predict subjects' excess path length on a given trial contingent on their past history of experience leading up to the current trial. Therefore, at the start of each delivery path we replaced the contents of the ideal navigator's memory with the structures that had actually been visible to the subject, from the start of the subject's session until the moment that was being simulated.

Clearly, in both experiments the ideal navigator greatly out-performs actual human subjects. Although the superiority of the ideal navigator was not unexpected, given the small size of the virtual towns that were being navigated the discrepancy between ideal and real navigators is surprisingly large. The magnitude of the discrepancy confirms that one or more of the ideal navigator's parameters substantially overstates the quality of the corresponding function in real navigators. To identify which functional modules(s) had been overstated in implementing the ideal navigator, we systematically degraded two of the ideal navigator's three modules (see Figure 4). The virtual environments were sufficiently compact, either 5×5 (for Dataset One) or 6×6 (for Dataset Two), and their layouts sufficiently simple, we made the assumption that actual navigators' ability to generate efficient paths—if they knew their own location and the location of a target store—was unlikely to differ much from the ideal navigator's ability to do the same. Therefore, we elected to retain the path generation module in non-degraded form as we used the ideal navigator as a benchmark against which to assess what impact imperfect memory (storage and recall) and/or imperfect vision (data acquisition) would have on human subjects' navigation through small virtual towns. Before turning to the results of degrading MAGELLAN's modules, note the difference in the behavior of the ideal, non-degraded navigator on the first delivery in Dataset One and Dataset Two (compare Figure 5's two panels). The ideal navigator's greater excess path length on Dataset Two's first delivery arises from the lack of a passenger search phase in the task that generated Dataset Two. In particular, because subjects began their first delivery with a passenger already on board, there was curtailed opportunity for the ideal navigator to build a spatial map before picking up the passenger. As a result, the navigator was forced to begin its initial delivery making use of a spatial map that was less complete than the spatial map available prior to the first delivery in Dataset One.

Degrading MAGELLAN's vision and memory

To identify possible sources of discrepancy between human and ideal navigators, we systematically degraded the ideal navigator's vision and memory, seeking to find the degradations that best brought the model into line with human navigators. This strategy resembles the approach introduced by Geisler (1989) and others in the sensory domain, and by Robles-De-La-Torre and Sekuler (2004) in the domain of precision movement control.

We systematically altered MAGELLAN's vision parameter, V, and its memory parameter, M. Variations in the V parameter governs which environment structures were added to the model's spatial memory; variations in M controlled how long structures in that spatial memory would remain viable. Thus, by adjusting the value of the V and M parameter, we control the degradation of MAGELLAN's *Vision* and *Memory* modules, respectively. Pseudocode for MAGELLAN is given in Appendix II.

Vision parameter. MAGELLAN's vision parameter, V, defines a threshold fraction of the screen, $0 \leqslant V \leqslant 1$,which a structure must occupy in order to be perceived. The V parameter is intended to account for failures to encode some structure displayed on screen. This failure could occur because the structure as displayed was too small, and/or because the subject brought insufficient attention to bear. Intuitively, the V parameter may be thought of as representing vision's spatiotemporal limitations (see, Geisler, 1989; Blake & Sekuler, 2005).[3]

[3] One set of circumstances required a modification of this application of the V parameter to be altered: If MAGELLAN were immediately adjacent to a structure, that structure was automatically added to the spatial memory –regardless of how much of the screen the structure occupied. This modification kept the model from becoming "stuck" when the V parameter was large, and the model's vision was so impaired that it could not see a building right in front

When V is set to 0, any structure that is displayed on screen is added to memory, and so MAGELLAN's vision component is identical to the vision component of the ideal navigator. As V increases, MAGELLAN's vision component grows less sensitive, and the cognitive map and route generation components are forced to operate with insufficient information. Finally, with $V = 1$ MAGELLAN's vision component is rendered entirely useless, and the now-sightless model must rely entirely on a blind search for a passenger. As MAGELLAN sees office buildings and stores, they are marked, and the cognitive map is filled in. MAGELLAN then tends toward unexplored blocks that are nearby its location, working outward toward unexplored blocks that are further away.

Memory parameter. MAGELLAN's memory parameter, M, is intended to account for forgetting over time; it governs the number of one-block steps after which a newly acquired memory will decrease to the point at which it is effectively no longer viable. The motivation for this parameter came from the observation that a subject may generate sub-ideal delivery paths even after having made one or even more deliveries to each store in an environment.

When $M = \infty$, once a structure has been added to the model's spatial memory it remains there permanently. As mentioned earlier, simulations of the ideal navigator incorporated $M = \infty$. In practice, MAGELLAN's memory behaves identically to the ideal navigator's, as long as M is greater than or equal to the total number of steps taken for all passenger searches (during which subjects find a passenger) and delivery paths (during which subjects deliver the passengers) in one environment. When $M = 0$, new information is not retained in spatial memory, and so purposeful navigation becomes impossible.

Fitting MAGELLAN to experimental data

To optimize MAGELLAN's account of subjects' empirical performance, we first minimized the root mean squared deviation (RMSD) between the mean *observed* learning curve and the mean learning curve *predicted* by MAGELLAN. Given that the model has only two free parameters we were able to find the best fitting values with a grid search of the two-dimensional parameter space, varying M in increments of 1 *step* and varying V in increments of 0.01 (equivalent to 1% of the screen). For Dataset One, the best-fitting set of parameters were $(M,V) = (36, 0.3)$, with an RMSD of 0.241 blocks. For Dataset Two, the best-fitting set of parameters were $(M,V) = (32, 0.07)$, with an RMSD of 0.692 blocks. Figure 6 shows the predicted and observed learning curves. Recall that MAGELLAN itself takes steps of one block each, which means that the model's predictions of performance in both experiments are accurate to within an average of less than one modeled step per delivery.

The quality of fit achieved to the mean learning curves suggests that MAGELLAN with sub-optimal memory and vision modules successfully reflects the average vision and memory degradation necessary to account for subjects' performance. However, the success of these fits does not guarantee that the model can also account for variability in difficulty across environments, or across trials within a given environment. To the extent that MAGELLAN approximates the way our subjects learned these environments, its performance should be able to tell us something about the relative difficulty of different environments. Difficulty of navigation in any environment is multiply determined. Difficulty reflects variability in the layout of stores and buildings; it reflects also where passengers are located, the sequence of their requested destinations, and the way the subjects' trajectory on previous deliveries influences the challenge posed by the current delivery. Recall that the experiments giving rise to the two datasets incorporated different levels of constraint in the way that stores were distributed within an environment. For Dataset One, stores were distributed in a quasi-random fashion, with minimal constraints. As a result, in some environments, stores happened to be clumped near one another, which should facilitate learning; in other environments, though, stores happened to be more widely disbursed, which could retard learning. In contrast, Dataset Two's stores were distributed in a more uniform fashion, which should make learning relatively uniform across different environments. Thus, if MAGELLAN is correctly predicting the difficulty of making deliveries in each environment, the model's performance should be highly predictive of the degree of difficulty of environments in Dataset One, but far less predictive of difficulty of environments in Dataset Two.

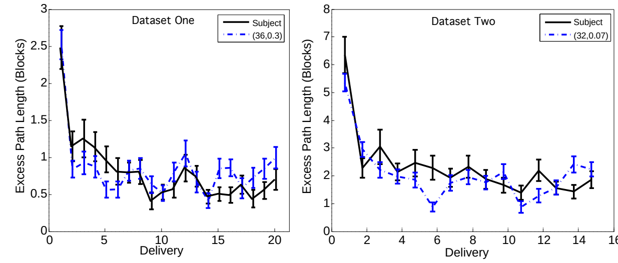

Figure 6. MAGELLAN's best fits to spatial learning curves from Dataset One (left panel) and Dataset Two (right panel). Error bars represent ±1SEM. For the model's fit to Dataset One, $(M,V) = (36, 0.3)$, RMSD = 0.241 blocks; for the model's fit to Dataset Two, $(M,V) = (32, 0.07)$, RMSD = 0.692 blocks. Recall that M is the number of steps taken before a new memory in the model's cognitive map becomes unusable. V determines a threshold fraction of the screen which a structure must occupy in order to be added to the spatial memory.

Because each subject learned a unique environment with unique passenger placements, it was not possible to fit MAGELLAN to the excess path length for each subject individually. Instead, we cast the data into contingency tables, which then were analyzed to verify the association between (a) parameters that fit the mean learning curves, and (b) the difficulty of learning different environments.

To analyze this association, we divided each dataset's results into three equally-populous bins, which were defined by

of it.

MAGELLAN's predicted excess path length for each subject. For this, we used the set of parameters that gave the lowest RMSD values for the mean learning curves. To generate the three bins, the model's predictions were sorted from shortest excess path length to longest. The predicted best third was grouped into the first bin, the second third was grouped into the second bin, and the worst third was grouped into the third bin. In this way, the bins correspond to predicted environment difficulty, with "easy" environments grouped in the first bin and "hard" in the third bin.

Next, we computed the corresponding *observed* excess path lengths for each of the three bins. For this we used only the first delivery to a given store, that is deliveries 1-4 in Dataset One, and deliveries 1-5 in Dataset Two. Because subjects learn to deliver efficiently after only a few deliveries, subsequent deliveries would not have provided sensitive indicators of environment difficulty. Figure 7 shows the outcome of the contingency analysis. For Dataset One, observed performance is monotonic with predictions based on the same parameter values that were used to make predictions in Figure 6. The reliability of the ordering was confirmed by paired t-tests, with $p < .0004$, $p < .00001$, and p=0.57, for bins 1 vs. 2, 1 vs. 3, and 2 vs. 3, respectively. However, as expected from the constraints used to place stores, the modest ordering seen for Dataset Two is not statistically reliable, with all three pairwise comparisons of bins yielding $p > .50$.

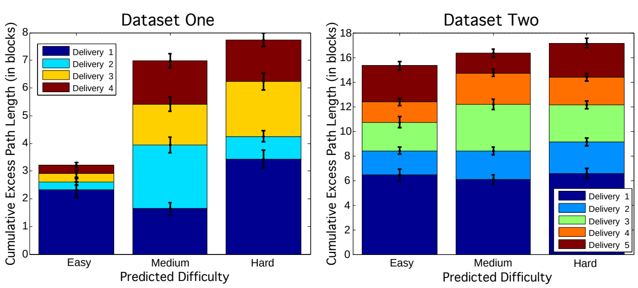

Figure 7. Predicted environment difficulty for Dataset One (left panel) and Dataset Two (right panel). Environments are binned according to MAGELLAN's predictions, using the set of parameters that best fit the mean learning curve for each dataset. Error bars represent ±1SEM. For Dataset One, the observed performance represented in each bin corresponds to the ordering of predicted environment difficulty; the ordering for Dataset Two is not statistically reliable.

For Dataset One, then, a single set of parameters fit both the mean learning curves to within the step size of the model, *and* correctly ordered environments by difficulty, as verified by observed subject performance in each environment. In particular, that subjects in predicted "easy" environments out-performed subjects in "medium" and "hard" environments suggests our simple model succeeds in capturing key, general trends in the empirical data. For Dataset Two, as expected from the uniformity of environments, observed differences among bins are negligible and not statistically reliable.

For Dataset One, the average number of steps taken between successive deliveries was 13.91. According to the best-fitting set of parameters for that dataset, memories for newly seen structures became unusable after thirty-six steps. Consequently, newly-acquired memories would have endured for an average of 2.59 deliveries. In Dataset Two, the average number of steps taken between successive deliveries was 11.45. According to the best-fitting set of parameters, memories for newly seen structures would have become nonviable after thirty-two steps. Consequently, memories here would have endured for an average of 2.79 deliveries. By assuming that memories become inaccessible between two and three deliveries after the corresponding structure was last seen, Magellan was able to fit subject's excess path length across deliveries in the two datasets.

General Discussion

We introduced MAGELLAN, a model of human spatial navigation based on the idea that subjects store associations between visible features in the environment and the spatial locations of those features. With just two free parameters, our model was able to account for mean performance in two virtual reality experiments. Using the same parameters, the model correctly ranked various environments in Dataset One as easy, medium, and hard to learn. This successful ordering is all the more remarkable for the relative homogeneity of the environments that had to be ranked. That is, all environments comprised the same number and types of structures, and were laid out on a rectilinear grid. The inclusion of additional constraints in the environment layouts used in Dataset Two minimized differences in predicted difficulty, making it impossible for MAGELLAN to correctly ordered the small differences among environments in that dataset.

Although MAGELLAN's psychological assumptions are surprisingly simple, its ability to perform a detailed, step by step analysis of subjects' navigation performance confers considerable power. By analyzing the information present on the screen during each of the navigator's steps, MAGELLAN was able to relate subjects' navigational accuracy to the information hypothesized to be available in the spatial memory.

Earlier, MAGELLAN's vision parameter, V, was defined in a physical metric as the fraction of the display that a structure had to occupy in order to be seen. That definition clearly understates the complexity of visual function. For one thing, selective attention plays an important modulatory role in the visual effectiveness of any stimulus (for example, Yeshurun & Carrasco, 2000; Carrasco & McElree, 2001). As a result, a future elaboration of MAGELLAN's V parameter might attempt to take account of not only stimulus properties but also attention-dependent effects. Of course such an elaboration would be useful only in the context of an experimental setting in which attention was varied independently of stimulus properties. One promising approach would be to exploit variation in gaze behavior in order to separate the V parameter's spatiotemporal limitations from attentional influences.

MAGELLAN as presently implemented has no representation of visual similarity, a variable that is known to exert considerable influence in various settings. As of now, the model assumes that any structure that is seen and entered into spatial memory will be distinct from any other structure that might be seen and entered into memory. As a result, MAGELLAN can retrieve the location of any stored structure without error, never confusing one previously seen structure with another. In the virtual reality experiments that generated Datasets One and Two, this assumption of completely distinctive structures may be justifiable. After all, subjects were able to give a verbal label to each store that was seen, and the store names in any single environment were designed to be highly distinctive. Additionally, prior to testing, subjects were familiarized with the name and appearance of all the stores that they might encounter. It seems likely, though, that in real-world wayfinding, people are influenced by the visual distinctiveness or differentiation of structures. City planners and researchers have long known that environments whose elements are differentiated, say in regard to architectural style or color, are easier to wayfind in than more homogenous environments. Particularly distinctive, highly-memorable structures enjoy a special designation – landmarks. Montello (2005) gives an excellent review of those characteristics of the physical environment that are known to influence orientation during navigation.

The virtual towns in Dataset Two were 40% larger in area than the virtual towns in Dataset One. Despite this difference, subjects still demonstrated an ability to produce near-ideal delivery paths after just a few deliveries. Studies of navigation in much larger towns, or with less distinct landmarks, should dramatically slow the rate of learning and allow for more detailed tests of the model's learning-related assumptions. The four-fold decrease in the best-fitting V-value from Dataset One to Two may be explained by the fact that Dataset One was generated with a lower resolution display than was Dataset Two (see Table 1). This lower resolution of the displayed environment could have forced subjects to see more of a store before being able to determine its identity, as fewer pixels – used to uniquely identify a texture-mapped structure – are packed into the same area of the screen. The need to view more of a store would translate, in turn, into a larger value of V. Further study is needed to explore this difference, including the possible impact of each subject's having to navigate six towns (Dataset Two) rather than just one (Dataset One).

Note that our treatment of an ideal navigator incorporated an assumption about the utility function that the navigator would attempt to minimize. Specifically, we have assumed that navigators seek to minimize distance traveled, an assumption that is consistent with the structure of our task—with subjects' bonus pay contingent upon excess path length. However, one can easily envision equally realistic driving tasks in which subjects attempted to minimize travel time rather than distance. Alternatively, subjects may wish to minimize travel along particular kinds of roads, such as toll roads, or they may plan routes that afford scenic views.

We should note that MAGELLAN's three modules have been treated here as independent of one another. In future work, that approach may prove to have been an oversimplification. After all, as Figure 4 illustrates, MAGELLAN's modules are interconnected, with several mutual dependencies. Working in concert, the modules' interactions produce a joint output—behavior, manifested in our task as the path taken while attempting to deliver a passenger. Interestingly, degrading any of the three modules will affect the shape and length of the path taken. Moreover, degrading each module could produce a distinct change in the delivery path's shape, as well as its length, which could affect another of the model's modules. For example, degradations in vision might alter the shape of paths around stores differently than around office buildings. These alterations in path shape could, in turn, influence the effectiveness of the navigation module. Additionally, degradations of memory could alter the path shape as buildings are dropped from memory, and so on.

Future extensions of MAGELLAN could generalize the memory module by allowing for variability in the number of steps before any particular memory is no longer usable. It seems reasonable that some structures are "forgotten" more quickly than others, perhaps as a function of depth of processing and/or as a function of structures' distinctiveness (Montello, 2005). The importance of perceptual differentiation in wayfinding can be likened to perceptual similarity's importance in recognition memory. Specifically, short-term visual recognition memory is influenced by two forms of similarity: inter-item similarity and/or target-item similarity (see, for example Kahana & Sekuler, 2002; Nosofsky & Kantner, 2006). Further on the subject of the memory module, we should note that we not committed to a forgetting process that is all-or-none. Rather, we could easily envision spatial location information being lost gradually as new, interfering information is learned. Similarly, information that has become inaccessible may later be retrieved in the presence of a salient contextual or associative cue (McGeoch, 1932; Tulving & Pearlstone, 1966; Wingfield, Lindfield, & Kahana, 1998).

The core idea behind MAGELLAN is that navigational learning relies on the storage and retrieval of associations between unique landmarks (stores or office buildings) and spatial locations. This notion of a cognitive map associating landmarks with spatial locations serves as the pillar of numerous recent models of spatial navigation (e.g., Byrne et al., in press; Howard, Fotedar, Datey, & Hasselmo, 2005; Leutgeb, Leutgeb, Moser, & Moser, 2005; Buzsáki, 2005; Gerstner & Abbott, 1997). As reviewed in the introduction, this idea derives considerable support from the well-known phenomenon of hippocampal place cells in rodents (O'Keefe & Dostrovsky, 1971), and the recent discovery of hippocampal place cells in humans (Ekstrom et al., 2003). Both humans and animals also possess cells that respond preferentially to distinct goals (Ekstrom et al., 2003; Hok et al., 2005).

MAGELLAN operates at a more macroscopic level than neurobiologically-oriented models of navigation whose focus is on the physiological coding of spatial information. Whereas those models concentrate on the details of spatial information's representation at both cellular and regional levels, our model assumes the existence of such neural codes, and takes off from that assumption. As an analytic tool for understanding navigation, MAGELLAN's strength lies in its use of the detailed behavioral data recorded in VR tasks to identify the information that is available to the system.

MAGELLAN shares its focus on naturally-occurring input with other current models in perception and cognition. For example, it has been known for some time that the statistics of natural images exhibit distinct regularities, not only in images' power spectra but also in their distributions of orientations (Olshausen & Field, 1996). Koch and Ullman (1985)'s saliency model suggests that regional variations in image statistics would afford information to guide eye fixations during visual search of complex scenes (Pomplun, 2006), including VR environments like those we have used. In particular, fixations in complex scenes specifically avoid areas that are uninformative, where informativeness is defined in terms of scene statistics, and, possibly, task-relevance (Kayser, Nielsen, & Logothetis, 2006). We believe that such theoretical constructs are highly relevant to possible future extensions of the work reported here. In its current implementation, MAGELLAN assumes that any object occupying a sufficiently large territory on the display will be seen and encoded. Thus, the model takes no notice of saliency of local visual features or of image content more generally. Arguably, though, all displayed regions of our virtual towns are not equally likely to be fixated, or entered into memory. It seems likely that any future implementation of MAGELLAN would gain in predictive power by incorporating parameters that reflect image-based differences in saliency (Peters, Iyer, Itti, & Koch, 2005), as well as possible task-dependent variations in looking behavior (Hayhoe, Bensinger, & Ballard, 1998; Pelz & Canosa, 2001).

A comprehensive model of spatial learning and spatial navigation is likely to account for learning of route-based information (landmark-to-landmark associations) alongside the learning of map-based information (spatial location to landmark associations), which is MAGELLAN's focus. The importance of route-based coding is clear from the considerable evidence for orientation dependence in human spatial cognition (e.g., Shelton & McNamara, 2001; Mou, Zhang, & McNamara, 2004). Benhamou, Bovet, and Poucet (1995) and Schölkopf and Mallot (1995) have both proposed route-based models of spatial information processing based on landmark-to-landmark rather than position-to-landmark associations. These ideas have a parallel in the literature on sequence learning, where Ladd and Woodworth (1911) first proposed that both position-to-item associations and chained item-to-item associations are important in serial learning (see also, Young, 1968; Lewandowsky & Murdock, 1989; Burgess & Hitch, 1999, 2005; Brown, Preece, & Hulme, 2000).

Although our very simple, two-parameter model gave a remarkably good account of performance in both datasets, what has been achieved is clearly only a first step toward a more complete account of spatial learning and navigation. Future studies should attempt to retard the rapid learning observed in our tasks. That would afford a more detailed analysis of the navigational learning curve. Future research might also examine the influences of navigation aids, such as global positioning systems (GPS), which lead drivers to desired locations. It would be valuable to determine, with the aid of the MAGELLAN model, how experience with such systems influences a driver's spatial memory for an environment navigated with the aid of GPS.

References

Benhamou, S., Bovet, P., & Poucet, B. (1995). A model for place navigation in mammals. *Journal of Theoretical Biology*, *173*, 163-178.

Bennett, P. J., & Pratt, J. (2001). The spatial distribution of inhibition of return. *Psychological Science*, *12*(1), 76-80.

Blake, R., & Sekuler, R. (2005). *Perception* (5 ed.). New York, NY: McGraw-Hill Companies, Inc.

Brown, G. D. A., Preece, T., & Hulme, C. (2000). Oscillator-based memory for serial order. *Psychological Review*, *107*(1), 127-181.

Burgess, N., & Hitch, G. (2005). Computational models of working memory: putting long-term memory into context. *Trends in Cognitive Science*, *9*(11), 535-41.

Burgess, N., & Hitch, G. J. (1999). Memory for serial order: A network model of the phonological loop and its timing. *Psychological Review*, *106*(3), 551-581.

Buzsáki, G. (2005). Theta rhythm of navigation: Link between path integration and landmark navigation, episodic and semantic memory. *Hippocampus*, *15*, 827–840.

Byrne, P., Becker, S., & Burgess, N. (in press). Remembering the past and imagining the future: A neural model of spatial memory and imagery. *Psychological Review*.

Caplan, J. B., Madsen, J. R., Schulze-Bonhage, A., Aschenbrenner-Scheibe, R., Newman, E. L., & Kahana, M. J. (2003). Human theta oscillations related to sensorimotor integration and spatial learning. *Journal of Neuroscience*, *23*, 4726–4736.

Carrasco, M., & McElree, B. (2001). Covert attention accelerates the rate of visual information processing. *Proceedings of the National Academy of Sciences, USA*, *98*(9), 5363-5367.

Ekstrom, A. D., Caplan, J., Ho, E., Shattuck, K., Fried, I., & Kahana, M. (2005). Human hippocampal theta activity during virtual navigation. *Hippocampus*, *15*, 881–889.

Ekstrom, A. D., Kahana, M. J., Caplan, J. B., Fields, T. A., Isham, E. A., Newman, E. L., & Fried, I. (2003). Cellular networks underlying human spatial navigation. *Nature*, *425*(6954), 184-8.

Geisler, W. S. (1989). Sequential ideal-observer analysis of visual discriminations. *Psychological Review*, *96*(2), 267-314.

Gerstner, W., & Abbott, L. F. (1997). Learning navigational maps through potentiation and modulation of hippocampal place cells. *Journal of Computational Neuroscience*, *4*, 79-94.

Graham, S. M., Joshi, A., & Pizlo, Z. (2000). The traveling salesman problem: a hierarchical model. *Memory & Cognition*, *28*(7), 1191-1204.

Hayhoe, M. M., Bensinger, D. G., & Ballard, D. H. (1998). Task constraints in visual working memory. *Vision Research*, *38*(1), 125-137.

Hok, V., Save, E., Lenck-Santini, P. P., & Poucet, B. (2005). Coding for spatial goals in the prelimbic/infralimbic area of the rat frontal cortex. *Proceedings of the National Academy of Sciences, USA*, *102*(12), 4602-4607.

Hollup, S. A., Molden, S., Donnett, J. G., Moser, M. B., & Moser, E. I. (2001). Accumulation of hippocampal place fields at the goal location in an annular watermaze task. *Journal of Neuroscience*, *21*(5), 1635-44.

Howard, M. W., Fotedar, M. S., Datey, A. V., & Hasselmo, M. E. (2005). The temporal context model in spatial navigation and relational learning: Toward a common explanation of medial temporal lobe function across domains. *Psychological Review*, *112*(1), 75–116.

Kahana, M. J. (2006). The cognitive correlates of human brain oscillations. *Journal of Neuroscience*, *26*(6), 1669–1672.

Kahana, M. J., & Sekuler, R. (2002). Recognizing spatial patterns: A noisy exemplar approach. *Vision Research*, *42*, 2177-2192.

Kayser, C., Nielsen, K. J., & Logothetis, N. K. (2006). Fixations in natural scenes: Interaction of image structure and image content. *Vision Research*, *46*(16), 2535-2545.

Keppel, G., Postman, L., & Zavortnik, B. (1968). Studies of learning to learn: VIII. the influence of massive amounts of training upon the learning and retention of paired-associate lists. *Journal of Verbal Learning and Verbal Behavior*, *7*, 790-796.

Kirschen, M. P., Kahana, M. J., Sekuler, R., & Burack, B. (2000). Optic flow helps humans learn to navigate through synthetic environments. *Perception*, *29*, 801-818.

Klein, R. M. (2000). Inhibition of return. *Trends in Cognitive Science*, *4*(4), 138-147.

Koch, C., & Ullman, S. (1985). Shifts in selective visual attention: Towards the underlying neural circuitry. *Human Neurobiology*, *4*(4), 219-227.

Korolev, I. O., Jacobs, J., Mollison, M. V., & Kahana, M. J. (2005). Human oscillatory activity during virtual navigation: A comparison between scalp and intracranial recordings. In *Abstract and itinerary planner, Society for Neuroscience.* Washington, D. C.: Program No. 65.16.2005.

Ladd, G. T., & Woodworth, R. S. (1911). *Elements of physiological psychology: A treatise of the activities and nature of the mind from the physical and experimental point of view (thoroughly revised and re-written.* New York, N Y: Charles Scribner's Sons.

Leutgeb, S., Leutgeb, J., Moser, M., & Moser, E. (2005). Place cells, spatial maps and the population code for memory. *Current Opinion in Neurobiology*, *15*, 738–746.

Lewandowsky, S., & Murdock, B. B. (1989). Memory for serial order. *Psychological Review*, *96*, 25-57.

Manning, J. (2006, May). *Modeling human spatial navigation using a degraded ideal navigator.* Senior Honors Thesis in Neuroscience, Brandeis University, Waltham MA. Waltham, MA.

McGeoch, J. A. (1932). Forgetting and the law of disuse. *Psychological Review*, *39*, 352-370.

McNaughton, B. L., Battaglia, F. P., Jensen, O., Moser, E. I., & Moser, M. B. (2006). Path integration and the neural basis of the 'cognitive map'. *Nature Reviews Neuroscience*, *7*(8), 663-678.

Montello, D. R. (2005). Navigation. In P. Shah & A. Miyake (Eds.), *The Cambridge Handbook of Visuospatial Thinking* (p. 257-294). Cambridge University Press.

Mou, W., Zhang, K., & McNamara, T. P. (2004). Frames of reference in spatial memories acquired from language. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *30*(1), 171-180.

Muller, R. U., Stead, M., & Pach, J. (1996). The hippocampus as a cognitive graph. *Journal of General Physiology*, *107*(6), 663-694.

Newman, E. L., Caplan, J. B., Kirschen, M. P., Korolev, I. O., Sekuler, R., & Kahana, M. J. (2006). Learning your way around town: How virtual taxicab drivers learn to use both layout and landmark information. *Cognition (in press).*

Nosofsky, R. M., & Kantner, J. (2006). Exemplar similarity, study list homogeneity, and short-term perceptual recognition. *Memory and Cognition*, *34*(1), 112-124.

O'Keefe, J. (1979). A review of the hippocampal place cells. *Progress in Neurobiology*, *13*(4), 419-39.

O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Research*, *34*, 171–175.

Olshausen, B. A., & Field, D. J. (1996). Natural image statistics and efficient coding*. *Network*, *7*(2), 333-339.

Pelz, J. B., & Canosa, R. (2001). Oculomotor behavior and perceptual strategies in complex tasks. *Vision Research*, *41*(25-26), 3587-96.

Peters, R. J., Iyer, A., Itti, L., & Koch, C. (2005). Components of bottom-up gaze allocation in natural images. *Vision Research*, *45*(18), 2397-2416.

Pomplun, M. (2006). Saccadic selectivity in complex visual search displays. *Vision Research*, *46*(12), 1886-1900.

Qi, D., & Sun, R. (2005). Learning to cooperate in solving the traveling salesman problem. *International Journal of Neural Systems*, *15*, 151-162.

Robertson, R. G., Rolls, E. T., Georges-Francois, P., & Panzeri, S. (1999). Head direction cells in the primate pre-subiculum. *Hippocampus*, *9*, 206-219.

Robles-De-La-Torre, G., & Sekuler, R. (2004). Numerically estimating internal models of dynamic virtual objects. *ACM Transactions on Applied Perception*, *1(2)*, 102-111.

Rolls, E. T., & O'Mara, S. M. (1995). View-responsive neurons in the primate hippocampal complex. *Hippocampus*, *5*(5), 409-24.

Samsonovich, A., & McNaughton, B. L. (1997). Path integration and cognitive mapping in a continuous attractor neural network model. *Journal of Neuroscience*, *17*(15), 5900-5920.

Sanchez-Vives, M. V., & Slater, M. (2005). From presence to consciousness through virtual reality. *Nature Reviews Neuroscience*, *6*(4), 332-9.

Schölkopf, B., & Mallot, H. A. (1995). View-based cognitive mapping and path planning. *Adaptive Behavior*, *3*, 311-348.

Shelton, A. L., & McNamara, T. P. (2001). Systems of spatial reference in human memory. *Cognitive Psychology*, *43*, 274-310.

Sholl, M. J. (1987). Cognitive maps as orienting schemata. *Journal of experimental Psychology: Memory, Learning & Cognition*, *13*(4), 615-628.

Taube, J. S., Muller, R. U., & Ranck, J. B. (1990). Head-direction cells recorded from the postsubiculum in freely moving rats. i. description and quantal analysis. *Journal of Neuroscience*, *10(2)*, 420-435.

Tolman, E. C. (1948). Cognitive maps in rats and men. *The Psychological Review*, *55(4)*, 189-208.

Tulving, E., & Pearlstone, Z. (1966). Availability versus accessibility of information in memory for words. *Journal of Verbal Learning and Verbal Behavior*, *5*, 381-391.

Vieira, F. C., Doria Neto, A. D., & Costa, J. A. F. (2003). An efficient approach to the travelling salesman problem using self-organizing maps. *International Journal of Neural Systems*, *13*, 59-66.

Voicu, H. (2003). Hierarchical cognitive maps. *Neural Networks*, *16*, 569-576.

Wingfield, A., Lindfield, K. C., & Kahana, M. J. (1998). Adult age differences in the temporal characteristics of category free recall. *Psychology and Aging*, *13*, 256-266.

Witmer, B. G., Baily, J. H., Knerr, B. W., & Parsons, K. C. (1996). Virtual spaces and real world places: Transfer of route knowledge. *International Journal of Human Computer Studies*, *45*.

Yeshurun, Y., & Carrasco, M. (2000). The locus of attentional effects in texture segmentation. *Nature Neuroscience*, *3*(6), 622-627.

Young, R. K. (1968). Serial learning. In T. R. Dixon & D. L. Horton (Eds.), *Verbal behavior and general behavior theory* (p. 122-148). Englewood Cliffs, N. J: Prentice-Hall.

Appendix I. Simplifying Subject Paths

In order to promote realism in the taxi driving task, subjects were allowed to drive freely on wide roads, making curved paths or zig-zag turns if desired. Their only restriction was the in-game collision detection engine, which prevented subjects from driving through buildings and walls or jumping up onto sidewalks and grassy areas. An important implication of this freedom of movement is that subjects can out-perform the ideal path as defined.

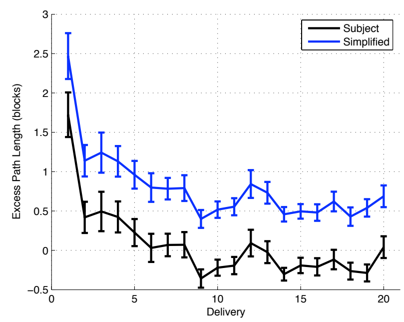

(a) Dataset One

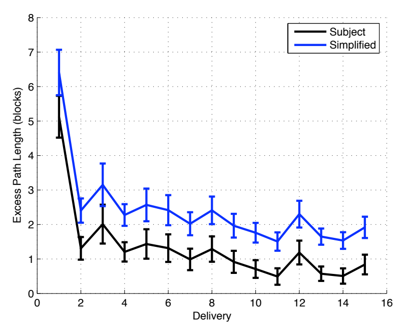

(b) Dataset Two

Figure 8. Excess path length vs. number of deliveries for Dataset One (a) and Dataset Two (b). Subject performance (black) is measured by comparing the average goal path distance to the average ideal path distance, for each delivery made. To compute the simplified subject excess path length (blue), we first round all positions to the nearest intersection. In this way, the simplified paths consist of straight lines and 90° turns. Note that while the actual excess path length can be negative, the simplified excess path length is non-negative. Error bars represent $\pm$1SEM.

Although the ideal path length is computed using the city block distance (plus adjustment for physical properties of the environment), the subject's delivery path has the ability to cut corners at headings other than 90°. Thus, the true ideal subject path is actually somewhere between the city block distance and the Euclidean distance from passenger to goal. In order to make a fair performance comparison between subjects and the model, we "simplified" the subject's paths to approximate the shape of a Magellan-style path as follows:

Because the VR towns are laid out on orthogonal grids, the heading along the simplifiedpath is always a multiple of 90°. Figure 8 demonstrates how path simplification affects the excess path length performance metric.

1. An $n \times m$ environment is divided into a $(3n+1) \times (3m+1)$ grid[4]of equal total dimensions. This means that if the $n \times m$ environment spans 256 virtual square *units*, the $(3n+1) \times (3m+1)$ grid will span 256 virtual square *units* as well.

2. The subject's path is plotted on the $(3n+1) \times (3m+1)$ grid.

3. Each point along the subject's path is rounded to the center of its block.

4. Rounded points are connected.

Appendix II. MAGELLAN pseudocode

```
function MAGELLAN(M,V):
  cogMap := {matrix of zeros, where each element e_(i,j) corresponds to the structure contained in the environment's i^th row and j^th column}

  for each delivery d do:
  subjForage := {subject's foraging path d}
forageSteps := {divide subjForage into block-sized steps}
for each step m in forageSteps do:
  for each structure seen at (x,y) during m (according to V) do:
  cogMap(x,y) := M
end
  cogMap(cogMap ≥ 1) := cogMap(cogMap ≥ 1) − 1
end
  temp := cogMap
arrivedAtGoal := False
driverPosition := POSITION({passenger d})
goalPosition := POSITION({goal d})
arrivedAtGoal := (driverPosition = goalPosition)
while not arrivedAtGoal do:
  if {temp({goalPosition}) > 0} then begin:
  {take ideal path from driverPosition to goalPosition}
break
end
  else begin:
  subGoal := {e_(i,j) ∈ temp = 0 nearest to {driverPosition}}
if {more than one subGoal} then subGoal := SELECT − RANDOMLY(subGoal)
driverPosition := TAKE − A − STEP(driverPosition, subGoal)
for each building seen at (x,y) during the step (according to V) do:
  temp(x,y) := M
end
  end
  arrivedAtGoal := (driverPosition = goalPosition)
end
  subjDeliv := {subjects's delivery path d}
delivSteps := {divide subjDeliv into block-sized steps}
for each step n in delivSteps do:
  for each structure seen at (x,y) during n (according to V) do:
  cogMap(x,y) := M
end
  cogMap(cogMap ≥ 1) := cogMap(cogMap ≥ 1) − 1
end
  end
  end
```

[4] Recall that each block is twice as wide as the road, so the total width of a road-block pair is three road-widths.